Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on initiatives, technologies, platforms, and products based on the Intel® Architecture.

Cover Story
Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles
Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor
To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives
Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue and is accessible beginning November 1999.

Bookmarking
We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

On behalf of all of us at Intel Developer Update, welcome to the future of the PC platform!

Table of Contents

(Click on page number to jump to articles)

COVER STORY **3**
THE NEW INTEL CHIPSET: ENHANCED GRAPHICS PERFORMANCE FOR THE PC PLATFORM 3

INSIDE LOOKING IN **5**
A CALLED SHOT 5

FROM THE EDITOR **7**

APPLIED COMPUTING **9**
DESIGNING A WEB-BASED MANAGEMENT SYSTEM THE RIGHT WAY 9

INITIATIVES & TECHNOLOGIES **13**
A STANDARDIZED APPROACH TO WEB-BASED ENTERPRISE MANAGEMENT 13
PLANNING FOR THE NEXT IA32 GENERATION 15

NETWORKING & COMMUNICATIONS **19**
DSL STANDARDS PROGRESS TO ENABLE MASS MARKET 19

SERVERS **21**
THE PC-AT BOOT PROCESS AND OPTION ROMS 21

WEB DEVELOPMENT **30**
SEYBOLD CASE STUDY AND TECHNICAL BRIEFS 30

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

The New Intel Chipset: Enhanced Graphics Performance for the PC Platform

Sandie Johnson
Product Marketing Manager
Platform Components Group
Intel Corporation

Overview

The introduction of two new Intel® Pentium® III processors at 600 MHz and 533 MHz coupled with the new, higher performance 133 MHz system bus, provides a wealth of new opportunities for hardware developers. But the advent of faster CPUs and system buses also requires comparable advances in chipset technology to guarantee a balanced platform with capabilities to meet a variety of needs.

Enter a new innovation from Intel designed to meet the challenge: the Intel® 810E chipset. Designed to accommodate the faster system bus while bringing new features to computer users, the Intel 810E chipset extends the capabilities of the Intel® 810 Chipset for Intel® Celeron™ processor-based PCs into the mainstream, while adding support for Pentium III processors and enhanced graphics performance.

The Bus Starts Here

The Intel 810E is the first system chipset to support a 133 MHz system bus, which raises the processor-to-memory and graphics bandwidth from today's top speed of 800 megabytes per second to a full one gigabyte/second. In addition, the new chipset also supports, 100 MHz and 66 MHz system buses, enabling it to provide developers with the flexibility to offer a single Pentium processor-based system board (motherboard) solution that spans mainstream and value price points in the PC marketplace.

The Intel 810E chipset incorporates Intel® Graphics Technology, building on the foundation laid by its 810 chipset predecessor. The Intel 810E also features the new Intel® Accelerated Hub Architecture, which at 266 MB/second provides twice the bandwidth of the PCI bus. This allows a much wider flow of content-rich application information from the I/O controller to its memory counterpart, providing optimized arbitration rules that enable more functions to run concurrently and pave the way for more lifelike audio and video on the PC platform.

The Intel Accelerated Hub Architecture includes a Graphics Memory Controller Hub, an I/O Controller Hub, and a Firmware Hub. The Graphics Memory Controller Hub provides the interface for the CPU, memory, and integrated Intel Graphics Technology, while supporting up to 512 MB of memory. The I/O Controller Hub forms a direct connection from the graphics and memory to an integrated Audio-Codec 97 (AC97) controller, an ATA-66 drive controller, dual USB ports, and PCI add-in cards. And the Firmware Hub stores system and video BIOS, as well as featuring the Intel® Random Number Generator (RNG), which enables stronger encryption, digital signing, and security protocols.

Performance Benefits at Value Pricing

Building on the Intel 810 chipset's integrated Intel Graphics Technology, which enhances performance by supporting both 133 MHz and 100 MHz local display caches, the Intel 810E expands 2D and 3D graphics performance for systems based on the Intel Pentium III and Intel Celeron processors. And as with the Intel 810 chipset, the new Intel 810E utilizes 100 MHz SDRAM technology.

Because of its support for a local 4 MB 133 MHz display cache, the Intel 810E chipset provides developers a seven percent performance boost for 3D graphics applications compared with the top-end performance of the Intel 810 chipset. The Intel 810E chipset and Intel's newest Pentium III platforms provide a 3D graphics performance boost of as much as 30 percent compared with low-cost Intel 810L chipset-based platforms.

The new 810E chipset enables developers to build platforms that deliver optimal performance and enhanced value through a variety of system efficiencies, including a reduced need for costly peripherals. The AC97 controller delivers stereo-quality sound, while 56 KB/second modem technology requires only the addition of an AMR riser card. Another innovation is the use of hardware motion compensation technology, which allows DVD playback without a separate decoder card.

Summary

One of the big advantages of the new Intel 810E chipset is the stability it provides developers. Because it's based on the foundation laid by the Intel 810 chipset, there's no need to completely redesign system boards (motherboards) to incorporate the new 810E. And because Intel conducts extensive tests on both hardware and software by using today's popular operating systems and peripherals, developers need not invest in extensive quality and validation processes.

Equally important, system boards based on the Intel 810E chipset have, all on one CD, the drivers and software to optimize the performance of the system. Developers can also reduce inventory costs, due to the fact that the Intel 810E chipset requires fewer discrete components such as sound cards, video cards, and modems.

Overall, Intel's new chipset technology paves the way for new levels in system performance, flexibility and stability, making the Internet and its content-rich applications more accessible and dynamic for end users everywhere. For system developers, that translates into a greater ability to provide customers with more and better choices than ever, spanning all price and performance points on the mainstream PC platform.

More Info

- 810E Product Overview
- 810E Design Guide
- 810E White Paper
- 810 Product Overview
- Random Number Generator
- Intel Graphics Technology

Author Bio

Sandie Johnson is the Intel 810 and 810E chipset marketing manager for Intel Corporation's Platform Components Group. Sandie has worked at Intel for over six years, holding several positions in the Information Technology and Platform Chipset Group Divisions. Prior to Intel, Sandie has worked more than 10 years in product marketing for Digital Equipment Corporation. She holds both a Bachelor of Science in Marketing and Computer Science from San Diego State University.

Inside Looking In

A Called Shot

Tim Mostad
Senior Technical
Marketing Manger
Intel Corporation

Overview

As summer slides to a finish, I'm reminded of the spectacle of the American professional baseball playoffs. They begin in September, after players from each team have played 162 games and expended months of sweat, blood, agony, and ecstasy, leaving a few to advance and even fewer to ascend to the rank of legend.

Perhaps the best known of baseball heroes, is the immortal George Herman "Babe" Ruth, whose fame owes to his homerun-hitting prowess. In his lifetime, he hit 714 over the outfield fence. Equally important to Babe's legend was his flamboyant style. He once stood at the plate, right arm extended at the outfield fence thereby announcing his intention to knock one out of the park, which he did. Babe's feat became known as the "Called Shot."

Imagine what this would mean in a business environment—for any company to be so confident that it could fully announce its intentions to its competitors, effectively pointing to the bleachers, then take a huge swing at a "homerun" product or technology and whack it out of the park. If we could hit a called shot, would we be legends too?

Of course, Babe Ruth didn't become a legend by hitting any one ball over the fence. His homerun record stood for 50 years and it took Henry "Hank" Aaron 4,000 more at-bats to surpass it. Accordingly, a company would have to have an impressive number of successes in a relatively short period of time to approach legendary status. Legends, it seems, need both a bigger than life image *and* stellar long-term performance.

Arguably, Intel has had its share of successes in its short existence, but has Intel ever hit a called shot? One effort worthy of consideration could be the establishment of the PCI Local Bus Interface Specification as an industry standard. We stepped up to the plate in face of a proven competitor, VESA Local Bus, and changed an entire industry with a single swing. Few doubt that PCI was a home run. But was it a called shot?

I was on the marketing team at bat, and I recall that we made some fairly grandiose predictions. "PCI will be *the* PC standard for the rest of the decade" was one of the tamer ones. These were big words from folks who'd never before stood at the plate. Anyone can boast, of course, so the fact that we were proven right doesn't exactly qualify us as legends.

A true element of greatness is illustrated by baseball's special dynamic between a pitcher and batter—face to face, nothing to hide, and nothing more than the tools in their hands and the knowledge in their heads. It is clear which competitor wins each engagement and how. The great players never resort to corked bats or spitballs—they play the game straight up. Likewise, a company wanting to hit the called shot would need to lay their technology and reputation on the line, and rely on winning simply with the best products based on that technology. If PCI approached this level, then the argument might begin to be compelling.

We knew early on that it would take more than just one company to deliver the force to launch PCI into the next millennium. We were smart enough to do an excellent job of enrolling and really incorporating the industry into the "PCI team." And then we did lay our PCI technology on the line through open licensing of the critical intellectual property, asking nothing more than for others in the field to do the same. It enabled us to enter the batter's box with the proper tool, the support and cooperation of the industry, in our hands. Without this, it would have been like hitting one over the fence without a bat. That would certainly have been legendary, but in the entire history of baseball it has yet to happen, and it won't. As with baseball, you need to go into business equipped with the essentials to play the game.

Our willingness to be open made it easy to share plans, goals, and maybe even some of our so-called secrets. I am sure some people would cringe if they knew how much information we exchanged. Yet, in the end, Intel profited, the industry profited, and the end user profited. It's hard to imagine a better outcome.

Whether or not PCI was actually a called shot may be irrelevant. Sportscasters who watched that famous incident disagreed as to whether Babe called the shot or not. Nevertheless, the legend has endured. The act of being willing to really be open to face the competition boldly, with the right means at hand, enabled us to help PCI go beyond what we could ever accomplish hidden away in some lab somewhere.

Maybe the called shot is just a sweet dream, made ripe on ball fields in the dwindling sunlight of fall days. And maybe today's competitive landscape is just too tough for standing at the plate without a "secret weapon." But I hope not. I just turned 40 and sadly it dawned on me that I'll never be a professional baseball player. I know that my only shot at the record book lies here in the computer industry, gazing at that fence, gesturing that I'll rip the cover off of the next one to cross the plate. I know that to do it I have to be part of a winning team that includes as much of the industry as will join the game.

Author Biography

Tim Mostad says, "the majority of my 18 years at Intel have been spent in the pursuit of technical marketing nirvana." He is responsible for demos, white papers, plugfests, and technical training to support the adoption of new desktop technologies.

From the Editor

Donna Loveland
Managing Editor
Intel Developer Update Magazine
Intel Corporation

Overview

Welcome to *Intel Developer Update*'s first issue. What you're seeing has been two years in the making.

Since September 1997, when we launched our first monthly online newsletter—on hardware only—we've been talking with all kinds of developers. Some of you have been consistently reading *Platform Solutions News*; some of you have never heard of it. Over time, we've learned a lot about what the entire developer community wants from a publication. In response, we've made a lot of adjustments along the way.

With this year's redesign of Intel's company-wide Web site we've done an overhaul, taking the opportunity to:

- *Expand our content*—so you can find material in *all* areas, hardware, software, and Web.
- *Use improved online tools*—so you can access information faster and more efficiently.
- *Refine our charter*—so you can get *news with context* in an article format from IDU, and go to developer.intel.com whenever you need fast facts, news items, event details, and data.
- *Change our name*—so we're tightly aligned with your other prime sources of information about Intel® Architecture, namely the Intel Developer Forum (IDF) and developer.intel.com.

In short, we've turned the *Platform Solutions* newsletter into *Intel Developer Update Magazine*.

We're kicking things off with a cover story on the new Intel® 810e chipsets, that became the hottest topic at last month's Intel Developer Forum.

For your convenience, IDU's articles are organized in departments, like you'll find in major magazines—nine in all:

- *Applied computing*—about embedded technologies and products
- *Desktop*—on everything from Basic to Performance PCs
- *Initiatives & Technologies*—what's coming out of our Labs and going into the industry
- *Mobile*—products and technologies for handheld and other portable devices
- *Networking & Communications*—spanning convergence technologies
- *Servers*—products and services for server solutions
- *Software*—everything from optimization to application development
- *Web Development*—for designers working with an eye toward Intel® Architecture
- *Workstations*—products from entry-level to high-end for industries from DCC to research

We'll post these departments only when they contain material that's worthy of your attention—no placeholders allowed. The goal is to bring you up-to-date once a month on the freshest information about Intel Architecture.

Whether you've been reading its predecessor, *Platform Solutions News*, or you're visiting Intel's pages for the first time, I think you'll like what you see in *Intel Developer Update*.

In fact, I encourage you to register for a free subscription. As a subscriber, you'll get a monthly email notifying you when we publish a new issue plus email news flashes when Intel announces big news for developers.

So, welcome. Enjoy. And come back often.

Author Bio

Donna Loveland is the editor of Intel® Developer Update. She joined Intel's Platform Marketing group earlier this year as the editor of *Platform Solutions News*. Donna began her high-tech career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held technical and marketing positions related to leading-edge technologies in areas ranging from stereoscopic display to digital broadcast to scalable online content. She holds a BA degree in English from the University of Rochester and an MA in Expository Writing from the University of Iowa.

Applied Computing

Designing a Web-based Management System the Right Way

James Blaisdell
Chief Technology Officer
Rapid Logic, Inc.

Overview

The Web browser has emerged as a "universal remote" for configuration, management, monitoring, and control of embedded network devices, including the growing number of devices based on Intel® Architecture platforms.

While most Web-based management strategies focus on embedding an HTTP (HyperText Transfer Protocol) server in the network device, the real challenge is how to get, set, and present device data elements in a systematic and extensible fashion.

Adding an HTTP server is the easy part. Just the opposite is true of the data layer, the mechanism that links the device data to the HTML user interface. While this may appear relatively straightforward, there are long-term consequences if it is done improperly.

At Rapid Logic we have learned over the course of more than 200 customer meetings that most vendors originally implement the data layer in the straightforward, intuitive, and wrong way. A software backplane architecture simplifies Web-based device control and makes interface design much less time-consuming and costly.

Why Web-based Management?

Three factors are driving the demand for Web-based configuration and management of network devices:

- The emergence of the graphical, user-friendly, platform-independent, and *free* Web browser as the ubiquitous point of network access.
- The ongoing industry movement to open standards, with TCP/IP emerging as the de facto standard for network-ready devices.
- The growing shift in adoption of networking technologies from homogenous, technical IT-centric enterprises to heterogeneous, non-technical users.

For all of these reasons, Web-based management is a "must-have" feature that network device vendors should plan to deliver in their products. For evidence of this trend, you need only look as far as the 70+ member Web-based Enterprise Management Consortium (WBEM) led by Cisco, Hewlett-Packard, and Intel. These vendors are all shipping or already delivering Web-enabled products.

Buy vs. Build

With the growing demand for Web-based management, network vendors face the challenges of increasing competition, the proliferation of management interfaces, escalating support costs, scarce engineering talent, and tighter time-to-market schedules.

Historically, device vendors have developed hardware, proprietary software, real-time operating systems, and TCP/IP stacks from the ground-up. Faced with the new challenges, vendors will increasingly focus on in-house development of only the 15 to 25 percent of the product that represents their core competency.

The widespread use of Intel Architecture (IA) for network devices provides software developers with a development platform whose reference design incorporates the same system architecture as the target hardware. Simulation and debugging are easier, because these activities can be carried out on a PC or workstation. There is extensive software tools support, and a variety of IA compilers are available to support leading operating systems.

Web-enabling a Device

Two key components are necessary for Web-enabling an embedded device. First, the HTTP server must be embedded in the device itself, where HTTP is the protocol of the World Wide Web. Secondly, a software mechanism must be embedded to link the device's dynamic data to the Hyper Text Markup Language (HTML) pages that constitute its user interface (UI). We call this mechanism the "data layer."

At first glance, an embedded systems engineer might assume that the HTTP portion of a project is the more time-consuming to develop. In fact, the development of the HTTP server can prove time consuming, but not for the reasons most engineers expect.

Most programmers assume that HTTP is time-consuming to deal with because they expect it to be a really huge protocol. This a reasonable assumption, when one considers that most commercial Web servers are large multi-megabyte applications. Driven by such assumptions, most developers feeling time-to-market pressure readily port a freeware version of an HTTP server, such as the NCSA server or the Apache server, instead of developing their own.

HTTP Is the Easy Part

In reality, HTTP is a relatively straightforward protocol that consists of nothing more than a lightweight messaging scheme that organizes different types of data, such as graphics and text. The typical browser-server session is simply a series of HTTP messages passed back and forth between the two entities.

HTTP also is stateless. This means that an HTTP server simply passes a message to a browser solely because that browser made a request for the message, not because the server was keeping track of the "state" of the HTTP session. As such, the browser is responsible to ensure that it gets all relevant HTTP messages. And if a given message is not received, most browsers will just display the data that they have, resulting in the occasional Web page with broken graphics or incomplete HTML text.

Developing an HTTP server is not time-consuming. The abundance of embedded HTTP servers on the market with a small footprint (8 Kbytes) or less provides testimony to the ease with which a server can be developed.

Dealing with the Data Layer

Many embedded systems engineers, when thinking about Web-based management, assume that the data layer, the software mechanism for linking the device data to the HTML UI, is relatively easy, compared to the HTTP server.

The reason for this assumption is `fprintf`, an ANSI C command supported by virtually every compiler and operating system in existence. To link a sample embedded HTML string with dynamic device data with fprintf, the embedded system engineer only has to write the following code shard:

```
fprintf( http_port, "<HTML><BODY> The device name is: %s <BR> The device address is: %d.%d.%d.%d
</BODY></HTML>", deviceName, ipAddr[0], ipAddr[1], ipAddr[2], ipAddr[3] );
```

From there, the engineer can just repeat and modify this example until the device's Web-based management starts to take shape.

Problems with `fprintf`

There are three points to make about the `fprintf` approach to the data layer:

- It is a very straightforward method. When Rapid Logic did our first spin of Web-based management we used this model. At the time, we thought this would enable us to get the product out the door in the least amount of time.
- This model is used extensively in most Telnet implementations. And since Telnet is the most common means of managing a network device that has been devised to date, there is a lot of legacy available to learn from.
- Most embedded developers in the industry are either using the `fprintf` technique, or planning on using it for the two reasons listed above. Of our prospective customers who had a Web-based management solution, 95 percent used `fprintf`s.

Most importantly, *this model is a horribly wrong way of doing things*. Unfortunately, the embedded developer won't know that it's the wrong way until he or she has to go back and modify the HTML user interface.

Avoiding the Traps

Vendors of Web-based management systems typically implemented command-line (CLI) management systems for their products. For these vendors, the most "intuitive" step is often to embed HTML code for GUI directly into the underlying C source.

Unfortunately, this method leads directly into a trap. Any changes to the GUI require invasive changes to the embedded HTML and recompilation of the critical, proprietary C code. This in turn requires a new embedded image that has a ripple effect on testing, fulfillment, version management, and product support.

Another common solution is to integrate C code into the GUI's HTML. While this results in a less disastrous trap, it nonetheless severely handicaps the design of the user interface (UI), including the ability to make changes. For example, this makes it difficult to move simple page elements such as buttons from one page to another. It also makes it impossible to propagate presentation objects to other management schema such as Simple Network Management Protocol (SNMP), CLI, or Java without re-engineering the product.

Don't Hard-Code that UI

There are potential problems with deeply intertwining C and HTML. It puts the onus of the UI development on the embedded system engineer, and it hard-codes the device UI.

Hard-coding has several negative consequences.

By hard-coding the UI, the engineer is unable to do rapid UI prototyping. This, in turn, limits the value of the device UI to whatever the embedded engineer comes up with first.

Hard-coding the UI also limits localization opportunities, forcing engineers to be adept at both language translation and C programming. Vendors are compelled to maintain separate code lines for each regional market.

Hard-coding limits OEM and channel opportunities. It is difficult enough to maintain different Telnet code lines for OEMs. It is unthinkable to do it for Web-based management.

Problems like these occur in all device management models that fail to create a clean abstraction for the data layer. In the RapidControl* Backplane model, the data layer is completely abstracted from the underlying elements.

The RapidControl Backplane was designed to function as a bidirectional protocol-independent MIB (management element) for getting, setting, and presenting a device's data. The RapidControl Backplane is essentially a unified data access model that features management protocol independence, "drop-in" support for legacy SNMP stacks, and a complete separation of front-end presentation from back-end proprietary C code.

While using a backplane is the most efficient and reliable way to provide true Web-based management, a backplane can also be approximated by using a proxy-based system. This has been called a "distributed backplane," even though it is a proxy system.

Although customer requirements tend to specify that Web-based management must be embedded in the device—which proxy systems do not achieve—in some cases a proxy system may be an acceptable implementation. If this is the case, the RapidControl Backplane can be easily configured to run in a proxy environment.

Faster Time-to-Market

RapidControl Backplane is protocol- and operating system-independent and helps to provide a flexible development platform for high-end network appliances such as email servers, fax servers, PBX systems, and Web servers based in the Intel Architecture.

The third key element of the ideal network development platform is the Linux* operating system (OS). Optimized for the Intel Architecture, Linux is an open-source OS that supports the technologies needed to develop scalable network appliances. The TCP/IP stack is built-in, and it provides both routing and dial-up support. Other advantages include simplified support for the easy addition, removal, and reconfiguration of software components.

With the combination of the Intel Architecture hardware platform, the Linux OS and the RapidControl Backplane, the network device designer's work is 80 percent done. This allows developers to focus their resources on the fast time-to-market development and put their effort into optimizing code.

Summary

Web-based management (WBM) strategies need to be evaluated according to three fundamental criteria:

- Ease of initial integration and time-to-market.
- Ease of future iterations and the ability to customize the Web-based graphical user interface (GUI).
- Life cycle extensibility to legacy device management technologies, such as Simple Network Management Protocol (SNMP) and command-line interfaces, and to evolving new ones, such as Java* and eXtensible Markup Language (XML).

The long-term impacts of a poor architecture for Web-based management are more serious than most other common errors. In short, the data-to-UI mechanism is the most fundamental piece when it comes to adding Web-based management to a network device. The Rapid Control software backplane represents an easy to integrate, easy to customize, extensible solution that simplifies the development of Web-based management systems.

Author Bio

James Blaisdell is a pioneer in Web-based management. As co-founder and Chief Technology Officer of Rapid Logic, Inc., he has worked extensively with many leading device manufactures to define their Web-management offerings, including Nortel Networks, PairGain Technologies, Tut Systems, RedBack Networks, Brocade Communications, and Transmedia Communications. He is the lead device engineer for Rapid Logic's MIBway for RapidControl, RapidControl for Web products, and the original architect for the RapidControl Backplane. Currently, he is involved with Rapid Logic's enterprise offerings and the lead device engineer for the approaching RapidControl for Network Appliances.

Initiatives & Technologies

A Standardized Approach to Web-Based Enterprise Management

Matthew Rollender
Strategic Marketing Manager
Network Interface Division
Intel Corporation

Vikas Aditya
Technology Assessment Manager
Network Interface Division
Intel Corporation

Overview

New applications such as Voice over IP (VoIP) and e-commerce have attracted a great deal of attention lately. In order to deploy these applications, the network must provide services such as QoS (Quality of Service) and Security (such as the IPSec protocol). To deploy these network-wide services, and also to enable deployment of future technologies, multiple network interfaces from different vendors should be managed in a consistent manner. One of the main problems currently facing developers is the complex variety of proprietary applications now available for collecting management information from network interfaces. This has made the provisioning of network services using today's management tools not only inefficient, but in many cases altogether impossible.

The Network Interface Services (NIS) initiative takes a badly needed standardized approach to solving this network management dilemma. Its objective is to effectively manage the complexity of end station network interfaces while keeping the total cost of ownership of the platform low. Initially developed by Intel and Microsoft and now being evaluated by the Distributed Management Task Force (DMTF), the NIS Initiative proposes a standardized means for sharing management information network interfaces located at the end stations in the enterprise via a standardized data model.

Already being employed as a standard, the Common Information Model (CIM) is an extensible data model of the enterprise-computing environment. Using CIM to exchange management information enables multi-vendor interoperability and allows management applications to focus on providing advanced network services. The NIS Initiative builds on the Common Information Model to help new technologies that are being incorporated into the end station network interfaces in today's enterprise networks. The data model proposed under NIS provides for representation of QoS and IPSec services provided by the network interface.

The Importance of WBEM

Developed by an industry consortium including Intel, Cisco, Microsoft, Compaq, and BMC Software, the Web-based Enterprise Management (WBEM) standard is currently backed by most major vendors. It provides three major benefits:

- Web-based access to management data
- Web-based end-to-end access to managed elements
- Data integration of management applications

WBEM allows existing access mechanisms such as the Simple Network Management Protocol (SNMP) and the Desktop Management Interface (DMI) to be used for legacy devices, while enabling management applications to share data in a consistent way.

The WBEM initiative has two main components. The first is the CIM extensible data description. CIM is intended to allow data from a variety of sources to be described and accessed, regardless of the original data source. It is both a management and instrumentation standard, which means that new devices can be instrumented using CIM directly.

The second WBEM component is standardized data encoding and transport. CIM can support a variety of distributed access methods, such as CORBA/IIOP and COM/DCOM. In addition, cross-platform remote access to CIM data is enabled through the newly released XML (Extensible Markup Language) encoding of CIM. XML encoding makes it possible to access the data using simple Web transport protocols such as HTTP. This in turn allows management solutions to be platform-independent and distributed across the enterprise.

NIS aims to standardize the data model for device configuration and monitoring aspect of the conceptual management model. NIS uses the CIM standard to model the functionality provided by network interfaces. This model is represented in the form of a schema, which is very similar to a UML diagram used in object-oriented design and modeling. This common data model allows different management applications to perform device configuration and monitoring on devices provided by different vendors. NIS doesn't replace existing management protocols; instead it provides an integration point through which data from all such sources can be accessed.

End-to-end Network Management

The successful implementation of network services means that many network interfaces must be manageable in a consistent manner so network devices behave and perform according to requirements defined by the services. Moreover, the use of Service Level Agreements (SLA) requires that services are provisioned and monitored to ensure that they conform to the agreements. NIS enables end-to-end network management that enables all applications and network devices to work together as objects within a single manageable system.

One of the principal benefits of the NIS initiative is that it frees management applications to focus on building network services, rather than worrying about how to exchange management data from proprietary end station network elements. NIS extends CIM schemas, or data models, for services provided by network interfaces, making the job of management applications much easier. Because it is based on industry standards being developed by DMTF, NIS means that management applications can inter-operate with equipment from different networking vendors and can also readily access advanced new features in the network interface in an open, standard manner. The end result is that this initiative will help shorten development times and enable application developers to bring more differentiated products to market faster.

The provision of a common interface for access to management data provides a common base that simplifies the task of developing network services applications. At the same time, a standardized approach makes it easier for IT managers to select solutions based on their merits.

More Info

A new Network Interface Services (NIS) white paper is now available from Intel.
WBEM and CIM specifications are available for download from the DMTF Web site.

Author Bio

Matthew Rollender leads a team in Intel’s Network Interface Division within Network Communications Group. His team is investigating new product requirements for Intel’s network adapter product line, and is working to drive protected, managed LAN communications product deployment for the enterprise. This work supports Intel’s initiative to make the trusted connected PC the platform of choice for managed communications on the Internet, as well as Intranet and extranet networking environments.

Vikas Aditya is a manager of Technology Assessment and Development team in Intel's Network Interface Division, which provides expansive line of server, desktop, mobile adapters as well as LAN on Motherboard products. His team investigates new networking technologies and standards on horizon and drives them into products.

Mr. Aditya holds MS degrees in Electrical and Computer Engineering from University of Massachusetts at Amherst and a Bachelor degree in Electronics and Communications from Delhi Institute of Technology.

Planning for the Next IA32 Generation

Bill Colson
Marketing Manager
Intel Architecture Labs
Intel Corporation

Overview

The era of the Internet is driving the desktop PC to the threshold of a revolution. Rich media and new Web-centric applications are literally adding a new dimension to the human interface, while making the computing experience more appealing for both sophisticated users and novices alike. In the next couple of years we will see a variety of specific, high capacity desktop computers that will embrace new broadband, speech, and vision technologies. Intel co-founder, Dr. Gordon Moore, made his first observation about the "doubling of transistor density on a manufactured die every year" in 1965. His statement, vindicated year after year, came to be known as "Moore's Law." In adhering to Moore's Law, Intel® Architecture continues to deliver the new processing power and system performance needed to deliver the rich content while keeping pace with new Internet applications

While all this sounds great from a marketing standpoint, actually putting Moore's Law to work can be a real challenge for developers. It means rolling up your sleeves and delivering differentiated PCs that meet the pricing and time-to-market constraints imposed.

As its name implies, Intel Architecture Lab's Scalable Platforms initiative provides you with practical ways to keep real-world platform capabilities in-step with the latest processor roadmap so elegantly described by Gordon Moore. At the Intel Architecture Labs, we take inspiration from another Intel founder, Dr. Robert Noyce, who advised "Don't be encumbered by past history, go out and create something wonderful."

We're doing it, in scalable ways.

Keeping up with the Roadmap

Advances in Intel® processor technology ensure that the next generation Internet will continue to run on Intel Architecture. To help developers stay ahead of the technology curve, the Scalable Platform initiative within the Intel Architecture Labs focuses on four goals:

- Removing platform technology roadblocks to the practical implementation of Intel's processor roadmap
- Optimizing the platform for both business and consumer applications
- Accelerating the performance of existing business and consumer applications
- Making the platform easier to use

Removing the Roadblocks

Intel focuses on supporting future processors through improved thermal performance in existing PC chassis designs, enhancing electrical power supplies, and by providing developers with innovative ways to drive down overall system costs.

Here is a brief summary of some of the latest platform technologies developed by Intel Architecture Labs:

NPSA

New Power Supply Architecture (NPSA) offers clear advantages over legacy power supply technology. Intel Architecture Labs is working to achieve the following design goals:

- An overall cost reduction in the area of 30 percent—50 percent cost savings compared to the existing system power supply box and 20 percent savings over present voltage regulator (VRM) design.
- Up to 80 percent efficiency—compared to 65 percent efficiency for conventional PS/2 power supplies, while maintaining the same form factor and size of existing units.
- Improved reliability with fewer parts—up to 20 percent improvement in usage of form factor area.
- Power Factor Correction (PFC) included at no extra cost.
- Cooling by the integrated 120mm-system fan—with no internal fan required in the power supply unit.
- Faster transient response—almost instantaneous voltage rise-time without “ringing.”

Board Manufacturing

Improvements in board manufacturing can trim the cost of the typical baseboard by as much as 10 dollars (US), while delivering improved EMI and thermal characteristics. Via-in-Pad (VIP) and Planes-on-Outer-Layer (POOL), a high density interconnect technology on baseboards, add-in cards, and memory modules can offer many advantages in size reduction, layer reduction, and increased routing capacity. By utilizing conventional through hole technology, via fabrication including size reduction and increased routing capacity can be achieved while still allowing for lower fabrication costs. VIP is required to make POOL work since VIP provides the necessary space for power/ground routing on the outer layers. VIP/POOL could deliver the advantages of an 8-layer ATX baseboard design (i.e., Reduced EMI) in a 4-layer PCB potentially saving 5 to 10 dollars (US) on each PCB. Baseboards with 10 layers can be reduced to 6-layer boards or less.

HAR Heatsink

High Aspect Ratio (HAR) heatsinks with heat pipes represent the migration of advanced mobile PC thermal management technology to the desktop PC. The technology allows a cost savings up to 20 dollars (US) per PC and the scalability to support processors up to 1.3 GHz, with no changes in board real estate requirements.

Fan Duct Cooling

This new way of mounting the PC’s cooling fan pulls air from the back of the chassis and ducts it downward onto the processor, chipset, memory, and graphics components. Without extensive retooling of the ATX or microATX chassis, fan duct cooling technology can effectively cool a 6-inch square board area. This thermal solution can essentially eliminate the need for expensive fansinks even permitting the cost-effective removal of some heatsinks from the memory or graphics components required on system boards today.

Optimizing the Platform

USB 2.0

At the Intel Developer Forum in February, Intel announced it was working on a version of the Universal Serial Bus that could run faster than 200-Mbits/second.

Then, at the Intel Developer Forum Fall'99, a revised target of between 30 to 40 times faster than USB 1.1 was announced. USB 2.0 data transfers are now targeted to have a maximum rate of 360-Mbits to 480-Mbits/second, based on engineering studies and test silicon by members of the USB Implementers Forum. Intel's partners on the technology include Compaq Computer, Hewlett-Packard, Lucent Technologies, Microsoft, NEC, and Philips. USB 2.0 will preserve complete backward compatibility with existing USB products, using the same cables, connectors, and software interfaces. The new USB 2.0 specification will appear in PCs in second half of 2000.

Intel expects both 1394 and USB to have important roles in the future of desktop and mobile computing, even as the roles for the interface are evolving. Strongly committed to the 1394 high-speed serial bus as the recommended connection between a PC and new digital consumer-electronics equipment, Intel considers 1394 a digital convergence pipe connecting the PC to the world of digital camcorders, digital VHS, set-top boxes, digital TV, and the like. The company believes convergence connectivity will grow increasingly important to the PC over time as the wealth of available 1394 devices increase. USB 2.0 is recommended for connecting PCs to computer peripherals, and 1394 is recommended as the interface connecting PCs to emerging digital A/V consumer equipment.

Future ATA

At the Intel Developer Forum in February, Intel announced it was working on a future extension of the ATA-66 interface that could act as a Gbit/s serial link. ATA is an excellent interface for internal storage devices, while USB is targeted for external devices. The vision of the serialized ATA derivative is to have greater cabling flexibility and greater throughput, while being backward compatible with current ATA-based hard disk drive devices.

Intel is working with multiple companies on the technology for a narrow, high-performance specification. For internal storage devices, the current ATA-33 specification, which permits a peak throughput of 33 Mbytes/s, will be replaced by the ATA-66 Gbit/s serial link specification this year. From 2000 to 2005, Intel and other industry leaders plan to introduce faster versions of the ATA interface with fewer signal pins. ATA serves an important role in today's PCs, allowing low-cost host controllers and low-cost disk drives with acceptable performance.

IALA

Intel® Application Launch Accelerator (IALA), dramatically speeds the time it takes for software applications to load from the hard disk by improving the performance of the personal computer’s storage subsystem. IALA is an optimized software algorithm for improved mapping to disk that lets applications launch much more quickly. The IALA technology actually minimizes the amount of travel the drive heads perform across the hard disk drive.

With IALA end users will see dramatic reductions in loading time when they launch applications on their personal computers. Running under Windows* 98, IALA can launch the Microsoft* Office* application suite two to three times faster. An IALA-enabled personal computer can launch all three Microsoft Office (Word*, Excel*, PowerPoint*) in approximately 3 seconds compared with the 12 seconds required to launch the same applications on a PC without IALA technology.

The technology is currently licensed to Microsoft for use in Windows 98 and Windows 2000. Developers can take advantage of the benefits of IALA technology by pre-optimizing the applications they load onto their system products to reflect the IALA. These user benefits can be realized without requiring any architectural changes to the overall personal computer, the I/O subsystem or the application itself.

User Friendly PCs

Connector Block

A one-piece, front connector block is being designed to locate USB, IEEE 1394 and RJ45 LAN, and modem connectors on the front of the PC for easier and more “appliance-like” user access. Placing required electrical components under the block liberates baseboard real estate.

IAPC

At Intel Developer Forum Spring ‘99, Intel (working with Microsoft and Toshiba) announced the industry's first total Suspend To RAM (STR) solution that fully supports Intel's Instantly Available PC (IAPC) initiative for reduced power consumption and rapid restart. With an Instantly Available PC, users can let their PCs enter a power saving "sleep state" during idle periods and avoid the time-consuming shutdown/restart/reboot process to interrupt and resume operation. When needed, Instantly Available PCs restore all applications and features to operation in a few seconds. When idle, STR-enabled systems consume only a small fraction of the power used for full operation.

Suspend To RAM is a cost-effective, optimal implementation of the Advanced Configuration and Power Interface (ACPI) 1.0 specification and Instantly Available PC design guidelines. ACPI brings together power management techniques within a PC whether in hardware, BIOS, operating system or application software. ACPI is a core part of Intel's mobile power initiative (MPI), Wired for Management (WfM) initiative, and Instantly Available PC technology.

Summary

The small/medium/large business professional has a vision—be Internet-centric, be scalable, be opportunistic, and be accurate. The end-user has a vision—be Internet connected, be Internet scalable, be Internet fast, and be Internet secure.

The number one device going into the 21st century to meet these visions is the personal computer. As a result, PC manufacturers are adjusting from being vendors of powerful and expensive boxes to becoming low-cost producers of highly efficient performance systems. Intel Architecture Labs is delivering the platform technologies that enable developers to design PCs with revolutionary processor and system performance to support the latest Internet applications, while also making next-generation PCs more cost effective, simpler to manufacture, and easier to use.

Intel Architecture Labs, established in 1991 and led by director Craig Kinnie, develops advanced technologies, such as videoconferencing, and creates industry standards upon which others build products. Unlike other corporate research labs, more than 90 percent of IAL's work concerning the Scalable Platforms initiative ends up in Intel products, as well as the wares of Intel's partners. Most other corporate labs report around 25 percent of research ending up in products.

If your job is platform implementation, take your cue from Robert Noyce, "Go out and create something wonderful."

More Info

For more details on how IAL's Scalable Platform initiative is removing technology roadblocks concerning IALA, Fan Duct, and PCB technologies, visit these Intel Web sites:

- Intel Developer
- The USB implementers forum
- Design/developer details
- Instantly Available technology
- IALA

Author Bio

Bill Colson is a 15-year Intel veteran. In his role as Intel Architecture Labs marketing manager, he is primarily responsible for platform architecture technologies. He holds a patent in the area of server management and is a member of the IEEE and Next Generation I/O Forum. Bill has written articles for EE Times and Electronic Design. He has presented at numerous industry forums and tradeshow events. Bill holds a BS in computer systems engineering and electrical engineering from the Oregon Institute of Technology.

Networking & Communications

DSL Standards Progress To Enable Mass Market

Andrew Hendry
Product Marketing Manager
Intel Architecture Labs
Intel Corporation

Overview

The gradual shift in the computing landscape toward affordable personal computers and dependence on the Internet has brought telecommunications to a significant inflection point. Over the past few years, personal computers have dropped in price and have become available worldwide. The Internet has become essential to the lives of millions of consumers, and its usage continues to grow substantially as consumers and businesses take advantage of Web-based applications and a wealth of online information. Consumer behavior has adapted to reap the benefits of the Internet by not only using it as a source of information, but also as a means of enhancing personal and business communications.

At the same time, consumers are demanding more from their Internet experiences as richer multimedia content and e-commerce capabilities become increasingly pervasive. To get the most from these applications, PC users require not only powerful PCs with high-performance processors, but also the high bandwidth transmission capabilities available with DSL (digital subscriber line), cable, and wireless broadband technologies. High-speed Internet connectivity for most consumers is not a question of if, but when. Faster Internet connections have not advanced as rapidly as desired for a multitude of reasons, resulting in the frustratingly slow Internet experience commonly known as the “World Wide Wait.” Consumers expect Internet performance that mirrors their PC experiences, and until recently, their needs have largely gone unmet.

High-Speed Access

In 1997, DSL service providers began touting high-speed access up to 27 times faster than existing dial-up access with a 56-Kbps analog modem and up to 12 times faster than ISDN access. Using the existing copper wire infrastructure (regular phone lines), DSL can be installed on the same line with voice service. This enables the user to simultaneously talk on the phone and perform high-speed data connections on a single line. The exponential performance improvement, use of existing infrastructure, and competitive, flat monthly rates for unlimited use makes DSL more attractive for Internet access and corporate LAN remote access than services such as ISDN and analog.

Emerging Market segments

The consumer DSL market segment is in a transitional period typical of emerging market segments. Although there are several types of DSL, the primary technology for the consumer mass market is Asymmetric DSL (ADSL). To meet significant consumer demand for high-speed Internet access, service providers have begun to offer proprietary DSL solutions. For the moment, since interoperable standards-based products are still under development, DSL providers are deploying these non-standard ADSL solutions, which are based on the ANSI T1.413 specification or Carrierless Amplitude Phase (CAP) technology, in “pairs.” This means the Central Office (CO) line card modem in the DSL Access Multiplexer (DSLAM) and the DSL modem at the customer premise must come from the same manufacturer. As a result, service providers must also supply these modems with the DSL service because consumers cannot obtain the proprietary DSL modems through traditional channels such as retail stores or PC OEMs.

ADSL Implementations

Until recently, service providers have deployed non-standard ADSL service with a splitter, a device that splits the voice and data frequency spectrum and requires an equipment installation by a service provider technician. These "truck roll" visits to customer premises are costly and do not scale to mass-market volumes. Now DSL providers are beginning to use proprietary splitterless ADSL implementations with distributed microfilters, thus making DSL equipment customer-installable. In the short term, this is a positive step because it simplifies DSL installation and hastens DSL service delivery to customers. In the next year, however, service providers will be able to deploy standards-based, interoperable and consumer-friendly DSL products that will enable greater availability and scalability for DSL.

Industry Standards

With Intel Architecture Labs (IAL) leadership, significant technical progress in the industry on standards, interoperability, and ease of use in the past year will enable a higher DSL adoption rate in the market segment. On June 22, 1999, the International Telecommunications Union (ITU) approved worldwide standards for ADSL: G.lite and G.dmt. G.lite is a consumer-friendly splitterless version of ADSL capable of providing data over existing phone lines at rates of up to 1.5-Mbps downstream and 512-Kbps upstream, and G.dmt is a splittered and splitterless version capable of providing data rates of up to 8-Mbps downstream and 2 Mbps upstream. In addition, the ITU also approved a new handshake protocol called G.hs. This collaborative effort by the industry leaders has created an environment in which the good of the overall industry is more important than that of individual equipment vendors.

Through the Universal ADSL Working Group (UAWG) and the ADSL Forum, IAL has also coordinated interoperability testing with equipment vendors and service providers for G.lite and G.dmt. This effort has involved creating interoperability test plans, coordinating all interoperability "plugfests," and organizing staging activities for public interoperability demonstrations.

For example, IAL led the planning for and implementation of the "G.lite Interoperability Showcase" at Supercomm '99, which was sponsored by the ADSL Forum and the UAWG. Thirty-three DSL vendors participated in this demonstration, which drove G.lite interoperability development and demonstrated that DSL is ready for the consumer mass market. In addition, IAL has conducted several G.lite field trials to verify and document the performance and ease of use of prototype G.lite service from ILECs and CLECs worldwide. At this point, the stage has been set for both ADSL service providers and vendors to benefit considerably from their contributions to standards and interoperability. Like the V.90 modem model, G.lite will lead the way for widely deployed standards-based DSL modems.

Summary

Until recently, the local phone companies were slumbering giants on DSL. But faced with tremendous customer demand and intense competition from high-speed cable companies and other DSL providers, these giants have awakened. And now, all are actively preparing for and implementing large-scale DSL rollouts to their customers encompassing technical, operational, and market segmenting activities. Standards-based, interoperable, and consumer-friendly DSL products will enable DSL providers to deploy service widely and scale with demand. As a result, by the end of 2000 nearly 2 million consumers and businesses are expected to be using DSL worldwide.

Author Bio

In his current position at Intel Corporation, Andrew Hendry is leading the DSL activities for the Connected.Home Initiative in the Intel Architecture Labs (IAL). The mission of the Connected.Home Initiative is to establish affordable and easy-to-use broadband network connectivity for consumers to and throughout the home. He has been involved in the Universal ADSL Working Group (UAWG) and the ADSL Forum, and works closely with industry partners to accelerate the adoption of high-speed communications standards. Mr. Hendry also facilitates trials of DSL technology between equipment vendors and service providers worldwide and works with internal engineering resources.

Mr. Hendry holds a MBA degree from the University of Michigan Business School. He also obtained a B.A. in Government from Cornell University.

Servers

The PC-AT Boot Process and Option ROMs

Roy Wade
Staff Software Engineer
Storage Components Division
LSI Logic Corporation

Ramin Neshati
Developer Guides
Technical Program Manager
Intel Corporation

Overview

The ubiquitous PC-AT architecture has reached a point where system designers and Independent Hardware Vendors (IHVs) are actively looking for ways to extend its capabilities, performance, ease-of-use, scalability, and other design characteristics. Much has changed since the introduction of the PC in the early 1980s: the system processors, I/O bus, memory map, boot process, graphics, etc., have all undergone drastic improvements. However, we stand today at the cusp of even greater technology transitions in systems design, allowing for vastly improved system performance, compatibility, multi-vendor interoperability, and "headroom" for scalability, future expansion, and growth.

The introduction of the IA-64 processor family allows system designers to make a clean break from the past and introduce new and innovative system designs that will do for server systems what the earlier Intel x86 processors did for the desktop PC—a massive proliferation of high-performance, low-cost server systems and solutions. Attempts at reducing or removing legacy technologies, allowing room for innovation, and extending the ease-of-use, performance, and scalability characteristics of server systems based on the Intel® Architecture have received special attention recently. A group of leading companies in the computer industry has formed working groups to collectively create an evolutionary path for the transition away from legacy technologies. LSI Logic and Intel, both leading companies in the markets they serve, have been at the forefront of paving the way for extending the Intel Architecture into the enterprise and the data center.

The PC-AT boot process lies at the core of this legacy migration trend. The introduction of the Extensible Firmware Interface (EFI), an abstraction interface that allows the de-coupling of the hardware from the operating system boot loader and provides other run-time services, has been enthusiastically adopted by the industry. However, the initial release of the EFI specification has not fully addressed an area that is of special importance to IHVs: option ROMs. However, efforts are already afoot to address this issue in the next release of the EFI specification. Quite simply, an option ROM allows the IHV to supply additional firmware to boot and configure its peripheral device and to extend the functionality of the system firmware (BIOS) during the boot process. There is no shortage of documentation detailing the PC-AT boot sequence. As PC technology has evolved, the number of specifications covering this topic has grown resulting in a fragmented image that takes some effort to grasp.

This article provides a comprehensive discussion of the PC-AT boot process in terms of today's technologies and architectures. Specifically discussed are PCI expansion ROMs for Plug and Play (PnP) devices and the details of their role in the boot process. A companion white paper that describes a legacy-reduced system boot process is currently in the early stages of formulation. We will defer the discussion of EFI option ROMs to this near-future effort.

Option ROMs Defined

The following is a general description of the current IA-32 option ROM architecture. Only Plug and Play compatible option ROMs are discussed. For a comprehensive and historical perspective, refer to the Plug and Play BIOS specification.

As was mentioned, an option ROM is a firmware component associated with a specific add-on device. The associated device may reside on the system-board or on a Plug and Play card. Accordingly, a device's option ROM will reside in system ROM or in PnP card ROM. Option ROMs are intended to isolate a hardware device by providing an abstract interface that implement device-specific functions, including:

- Power-on self-test
- Initialization
- Interrupt service routines
- BIOS routines

The I/O services provided by an option ROM serve as a translation layer between differing protocols. For example, an option ROM may provide a legacy Int 10H video interface to an AGP device; or provide a legacy Int 13h disk drive interface to one or more disk drives attached to a SCSI bus. The PnP option ROM model is generic and can support any Plug and Play device. Requirements specific to an application are detailed in existing specifications and are beyond the scope of this article (e.g., a PCI-to-USB PnP product would be concerned with the appropriate PCI and USB specifications).

Typically a PnP device will provide at least one option ROM image, but some architectures allow for more (e.g., PCI). It is not required that a PnP device provides an associated option ROM, hence they are optional. Option ROMs facilitate the Plug and Play concept by adhering to a common image format, which indicates flexible resource requirements in support of automatic configuration.

During the boot process, option ROM images are copied from their ROM to main memory (RAM) and their ROM address range is mapped (shadowed) to main memory. This is done to increase execution speed by taking advantage of the faster access time of RAM versus ROM. The system BIOS is responsible for this expansion process during boot. The system BIOS loads and initializes option ROM images at its discretion. In newer systems, only those option ROMs required to boot the operating system are loaded. The operating system then loads its own drivers for all PnP devices while shadowed option ROMs remain resident (and dormant) in RAM.

The primary purpose of the pre-boot subsystem is to provide a set of services to find and execute an operating system loader in a reliable and expeditious manner. The pre-boot environment can also provide error checking, error logging/reporting and error recovery as well as support system expansion.

Two chief components in the boot process are the system BIOS and Initial Program Load (IPL) device option ROMs. The system BIOS is a single instance entity that controls the sequence of boot events. It provides abstracted services, enumerates hardware devices, loads option ROMs, and performs resource management. The system BIOS gains control of the initial boot process by intercepting a power on or reset. It invokes the system bootstrap loader to load the operating system (possibly via an option ROM) from the IPL device and execute it. As discussed previously, option ROMs provide services to access the devices attached to their bus, in this case the boot device.

PC-AT Boot Model

The following discussion of the PC-AT boot model assumes that the reader is familiar with:

- PnP compatibility for all components (system BIOS, devices, option ROMs, operating system)
- PCI 2.1 compliance
- BIOS Boot Specification support for system BIOS and IPL device option ROMs
- Device Driver Initialization Model (DDIM) as specified in the PnP BIOS Specification

Figure 1 below shows an example system configuration that may prove helpful in a description of the boot sequence.

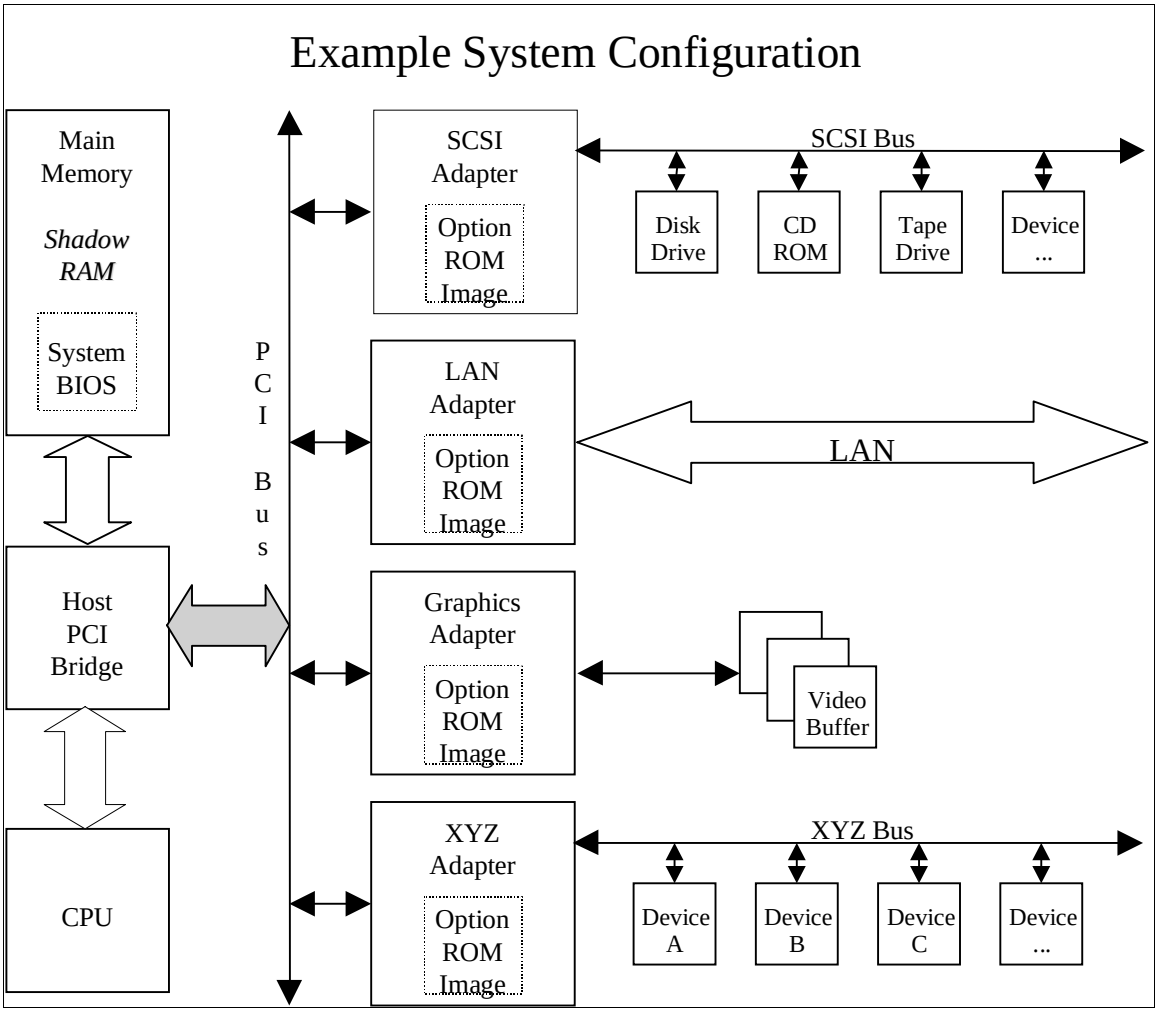

Figure 1 - Example System Configuration

The BIOS Boot Specification defines a method to specify the IPL devices to attempt to boot from and prioritize their order. This method requires that option ROMs identify their IPL devices to the system BIOS and that the system BIOS provides a mechanism for the user to control the IPL order. The comprehensive description of this mechanism can be found in the BIOS Boot (BBS) Specification and is not repeated here.

IPL Sequence of Events

Below is the sequence of events to define the IPL order. In this example, it is assumed that at least one IPL candidate device does exist.

IPL Sequence of Events:

- The system BIOS captures control of the system after power on or reset.
- The system BIOS performs the following functions:
- Initializes the primary output device (display).
- Initializes the primary input device (keyboard).
- Performs self-test functions (RAM test etc.).
- The system BIOS scans the system's busses (e.g., PCI) for PnP option ROM Headers.
- The system BIOS performs resource conflict resolution and allocation. The algorithm used in the resource allocation process is beyond the scope of this specification. In the case of PCI, the resource requirements (IRQs, I/O, memory, etc.) for an option ROM are specified in the Configuration Registers.
- The system BIOS displays a prompt that the operator uses to enter the Setup Utility.
- The system BIOS reads previous settings from non-volatile memory.

- The system BIOS scans the system's busses to detect PnP expansion headers for option ROMs that provide access to potential IPL devices. For each IPL device option ROM that is found:
- The system BIOS copies the ROM image into main memory.
- The system BIOS calls the image's initialization entry point, which is specified in the PnP option ROM Header. The option ROM performs the following:
- Performs required initialization, and determines the device(s) it is to control.
- If there are no devices to support, the option ROM "resizes" itself to zero in on its option ROM Header, then returns from its initialization entry point.
- If there are devices to support, the option ROM creates a linked list with a new PnP expansion header for each device found. The PnP expansion header fields will uniquely identify each device. A field of particular interest to the user will be the Product Name String. The option ROM then returns from its initialization entry point.
- At this point all option ROMs with PnP expansion headers have been initialized. Option ROMs with device(s) to support will not have hooked any interrupts (e.g., Int 13h).
- The system BIOS now scans main memory to detect all PnP expansion headers for potential IPL devices. For each PnP expansion headers that is found:
- If a BEV header, the entry is added to the IPL priority list.
- If a BCV header, the entry is added to the BCV priority list.

NOTE: See [BIOS Boot Specification Version 1.01] for the distinction between the IPL and BCV priority lists.

- The system BIOS presents the IPL and BCV priority lists to the user via a Setup menu.
- The user manipulates the IPL and BCV priority lists as desired, saves the settings, then exits Setup.
- The system BIOS saves the settings to non-volatile memory and automatically reboots the system.

PnP BBS System

We are now ready to discuss the sequence of events that occur during the boot process of a PnP BBS system. The focus is directed toward the interaction between the system BIOS and option ROMs.

Plug and Play Bios Boot System Sequence of Events:

- The system BIOS captures control of the system after power on or reset.
- The system BIOS performs the following functions:
- Initializes the primary output device (display).
- Initializes the primary input device (keyboard).
- Performs self-test functions (RAM test, etc.).
- The system BIOS scans the system's busses (e.g., PCI) for PnP option ROM headers.
- The system BIOS performs resource conflict resolution and allocation. The algorithm used in the resource allocation process is beyond the scope of this specification, see [PnP BIOS specification reference]. In the case of PCI, the resource requirements (IRQs, I/O, memory, etc.) for an option ROM are specified in the Configuration Registers.
- The system BIOS displays a prompt that the operator can use to enter the Setup utility. The prompt is ignored in this case.
- The system BIOS scans the systems' busses (e.g., PCI) to detect PnP expansion headers for option ROMs that provide access to potential IPL devices. For each IPL device option ROM that is found:
- The system BIOS copies the ROM image into main memory.
- The system BIOS calls the image's initialization entry point, which is specified in the PnP option ROM header. The option ROM performs the following:
- Initialization. It also determines if the device(s) it is to control.
- If there are no devices to support, the option ROM "resizes" itself to zero in its option ROM header then returns from its initialization entry point.
- If there are devices to support, the option ROM creates a linked list with a new PnP expansion header for each device found. The PnP expansion header fields will uniquely identify each device. The option ROM then returns from its initialization entry point.
- At this point, all option ROMs with PnP expansion headers have been initialized. Option ROMs with device(s) to support will not have hooked any interrupts (e.g., Int 13h).

- The system BIOS reads settings from non-volatile memory, including the IPL and BCV priority lists. If non-volatile memory is corrupt, default IPL and BCV priority lists are created.
- The system BIOS builds the run-time BCV priority list.
- Main memory is scanned to detect all PnP expansion headers that have BCV entry points.
- Each BCV entry that is found is appended to the run-time BCV priority list.
- The system BIOS compares the NVM and run-time BCV priority lists (actually it compares the number of items in the lists). For the purposes of this discussion we assume the lists are the same (if the lists do not match the system BIOS may clear the NVM BCV priority list). Either way, the run-time BCV priority list is used from this point on.
- The system BIOS calls each BCV entry (extracted from the PnP expansion header) in the order specified in the BCV priority list. For each entry in the BCV priority list, the system BIOS calls the associated BCV entry point to request the option ROM to install Int 13h services.

The option ROM performs the following:

- The option ROM determines if it actually controls any Int 13h IPL devices.
- If there are no IPL devices, the option ROM simply returns from its BCV entry point. If there are IPL devices, the option ROM proceeds.
- The option ROM determines the number of hard drives currently installed by reading BDA address 0040:0075.
- If no other hard drives are installed (i.e., BDA 0040:0075 is zero, so this is the first hard drive) the current Int13h vector is copied to the Int 40h vector so that floppy services are handled properly.
- The option ROM infers the drive number for the current BCV IPL device from BDA 0040:0075 (i.e., BDA 0040:0075 + 80h).
- The option ROM hooks the Int13h vector, saving the current vector in order to chain to services for other drive numbers. Note that the option ROM may install Int 13h services for more than one drive in a single “hook”.
- The option ROM increments BDA 0040:0075 by the number of drives it installed services for, typically the number of drives installed is one.
- The option ROM returns from its entry point.
- At this point the option ROMs for all IPL devices have been initialized, all BCV device Int 13h services have been installed, and shadow RAM has been read-only enabled.
- The system BIOS sets the Int 18h vector to the address of its failed boot attempt recovery entry. This function is responsible for walking through the IPL priority list on failed boot attempts.
- The system BIOS sets the Int 19h vector to the address of its boot strap loader. This function is responsible for loading and invoking an IPL device’s bootstrap loader or BEV entry point.
- The system BIOS executes Int 19h. The first entry in the IPL priority list is the specified boot device.

Int 19h Processing:

- If the current IPL priority list selection is a BCV device:
- The system BIOS reads the device’s bootstrap loader from Sector 1, Head 0, Cylinder 0 to address 0000:7C00. The Int13h services for the BCV device’s associated drive number is used to read the loader.
- If a valid boot sector is detected, the system BIOS loads the OS bootstrap loader and calls its entry point. If the OS loader fails, it executes an Int 18h to indicate the failed boot. If the OS loader is successful, the OS is in control and there is no return to Int 19h.
- If the current IPL priority list selection is a BEV device:
- The system BIOS calls the BEV entry (extracted from the PnP expansion header). The BEV entry may also hook Int 13h, for example if the BEV is for an “El Torito” CD-ROM Hard Drive emulation.
- The BEV loads and executes its OS bootstrap loader.
- If the BEV or OS loader fails, it executes an Int 18h to indicate the failed boot. If the OS loader is successful, the OS is in control and there is no return to Int 19h.
- Int 18h is executed to indicate a failed boot.
- Int 18h Processing
- If all devices in the IPL priority list have been attempted:

- An error message/prompt is displayed indicating no OS was found.
- When the user responds to the prompt, the system BIOS executes Int 19h. The first entry in the IPL priority list is the specified boot device (i.e., starts over again).
- Otherwise, the system BIOS executes Int 19h, specifying the next relative IPL priority list entry as the boot device.

Single Option ROM Code Image

We now direct our attention to the description of the components and format of a single option ROM code image. Option ROMs that follow the Device Driver Initialization Model (DDIM) (PnP BIOS specification reference) exist in two forms:

- Pre-initialization Image—the original option ROM image as programmed into the system-board or PnP card.
- Post Initialization Image—the option ROM image as it exists in shadow RAM after its initialization entry has successfully returned.

Option ROM Components

PnP Option ROM Header: This is the root element of a Plug and Play option ROM. It identifies an image as an option ROM and provides links to the PCI Data Structure and PnP expansion header. The system BIOS uses the information in this header to:

- Detect a PnP option ROM image
- Copy (shadow) an image into RAM
- Validate an image via a checksum
- Initialize an image by calling its initialization entry
- Locate an image's PCI Data Structure
- Locate an image's PnP expansion header

Start with BIOS Boot Specification Version 1.01 for a detailed description.

PCI Data Structure: This is a companion item to the PnP option ROM Header. It qualifies an image as a PCI Expansion ROM and contains identification information for an image and its device.

Refer to PCI BIOS Specification Revision for a detailed description.

PnP Expansion Headers: This is a companion item to the PnP option ROM Header. In a pre-initialization image it identifies an image as a PnP option ROM. In a post-initialization image, one or more of these headers may exist, each identifying an IPL device, including:

- Reference to Product Name String
- Reference to BEV Code
- Reference to BCV Code
- Reference to the next PnP expansion header (i.e., linked list of IPL devices)

Start with BIOS Boot Specification Version 1.01 for a detailed description.

Runtime Code: The Runtime Code is the constant data and executable code that remains after an option ROM has successfully returned from its initialization entry. This code includes:

- BEV Code (BIOS Boot Specification Version 1.01)—code to directly load an OS from an IPL device and optionally hook Int 13h.
- BCV Code (BIOS Boot Specification Version 1.01)—code to detect an IPL device and optionally hook Int 13h. Int 13h Service Code—conventional and enhanced [EDD specification reference] Int 13h disk drive services.

Initialization Code: The Initialization Code is the constant data and executable code that is used only during initialization then discarded before returning from initialization entry. This code is responsible for:

- Device detection
- Device initialization
- Initializing runtime data structures
- Creation of Product Name String for each detected device
- Creation of PnP expansion header for each detected device
- "Disposing" of itself per the DDIM, see (Initialization Sequence paragraph reference)

Pad: Unused space used to expand an image to the next greatest 2048 byte boundary. Legacy requirements dictate that an image's size be a multiple of 2048 bytes.

Checksum: The checksum of the entire option ROM image. A value such that the 8-bit sum, using unsigned 2's complement addition ignoring overflow of all bytes including the checksum, is zero.

Product Name Strings: A Product Name String is a null terminated ASCII string that is intended to uniquely identify an IPL device. Currently only the first 32 characters are displayed and therefore significant.

Start with [BBS reference] for a detailed description.

Initialization Sequence: The Plug and Play (PnP) BIOS Specification (PnP BIOS Specification Reference) introduced the Device Driver Initialization Model for option ROMs. In this model, an option ROM image exists in two forms:

- Pre-initialization option ROM Image (Figure 1)—the image as it exists in ROM on the system board or PnP Card (Runtime Code and Initialization Code).
- Post-initialization option ROM Image (Figure 2)—the image as it exists in Shadow RAM after returning from its initialization entry (Runtime Code only).

The DDIM defines the transition from the Pre-initialization to the Post-initialization option ROM image. This transition facilitates:

- More efficient use of Shadow RAM by allowing initialization code (and data) to be discarded
- A means of dynamically building data structures and saving them as static (i.e., constant) data at boot time

Option ROM Initialization Sequence of Events:

- The system BIOS copies the ROM image into main memory.
- The system BIOS calls the image's initialization entry point, which is specified in the PnP option ROM Header. The option ROM performs the following:
- Initializes common data items.
- Searches for device(s) to control, for each device found:
- Initializes the device (reset, etc.) and its associated data structures.
- Creates a Product Name String to identify the device.
- Creates a PnP expansion header for the device and links it to the PnP expansion header List.
- If there are no devices to support:
- Writes zero to the Length field in its PnP option ROM Header.
- Returns from its initialization entry point with a status code indicating "no devices to support".
- Or optionally, if the option ROM decides to leave a static footprint (An option ROM may want to leave an image resident in Shadow RAM to provide data to a related option ROM or driver, for example NVS data or Vital Product Data):
- Copies any data it wishes to be static to the end of the runtime image (i.e., overwrites code).

- Updates the Length field (Length = Pre-initialization Length – Code Length + Static Data Length) in its PnP option ROM Header. This effectively discards the Code and reduces the image's footprint.

- Calculates the new image checksum and appends it to the new image.
- Returns from its initialization entry point with an appropriate status code.
- If there are devices to support, the option ROM will arrange and reduce its image as follows:
- Copies any data it wishes to be static to the end of the runtime image (i.e., overwrites Initialization Code). This would include the Product Name Strings and the PnP option ROM Headers.
- Updates the Length field (Length = Pre-initialization Length – Initialization Code Length + Static Data Length) in its PnP option ROM Header. This effectively discards the Initialization Code and reduces the image's footprint.
- Calculates the new image checksum and appends it to the new image.
- Returns from its initialization entry point with a status code indicating "successful initialization".

This article has attempted to document the legacy PC-AT boot process with special emphasis on the role of option ROMs. As part of the on-going industry effort to migrate away from legacy technologies and architectures, Intel intends to publish additional information that will describe a legacy-reduced IA server boot environment.

More Info

The following documents provide details on option ROMs and the PC-AT boot process, some of which are referenced in this article:

- BIOS Boot Specification Version 1.01, Compaq Computer Corporation, Phoenix Technologies Ltd., Intel Corporation, 1996.
- BIOS Enhanced Disk Drive Specification Version 3.0, Phoenix Technologies Ltd., 1998.
- Clarification to Plug and Play BIOS Specification Version 1.0.
- "El Torito" Bootable CD-ROM Format Specification Version 1.0, Phoenix Technologies, Ltd., IBM Corporation, 1994.
- PCI BIOS Specification Revision (latest), PCI Special Interest Group, Hillsboro, OR.
- Plug and Play BIOS Specification, Version 1.0A, Compaq Computer Corporation, Phoenix Technologies, Ltd., Intel Corporation, 1994, or ftp://download.intel.com/ial/wfm/bio10a.pdf.
- POST Memory Manager Specification Version 1.01, Phoenix Technologies Ltd., Intel Corporation, 1997.

Author Bio

Roy Wade is a staff software engineer for LSI Logic Corporation. He is responsible for the design and development of BIOS and OS software to support LSI Logic's IO products. Roy has over 13 years of industry experience in a variety of areas, including satellite navigation, telecommunications, and Client/Server applications. Roy is a Phi Beta Kappa graduate of the University of North Dakota where he received his BS in Computer Science in 1986. Roy is a contributing member of the UNIX Developer's Interface Guide for IA-64 based servers (UDIG) working group. He has collaborated on patent applications in both the wireless communications and PC BIOS areas.

Ramin Neshati joined Intel in January 1998. Before becoming the developer guides technical program manager, he represented ESG on the Wired for Management initiative. Aside from Intel, Ramin has over 17 years of industry experience at S3 Incorporated, Dell Computer Corporation, Xerox Corporation and Link Systems Incorporated, in a variety of positions from software engineering to software architecture and engineering management. Ramin received his MBA from Pepperdine University (1993), the Master's degree in Computer Science from University of Idaho (1982) and the Bachelor's degree in Computer Science from Washington State University (1980). He has collaborated on several patent applications, published articles on networking protocols and services, and lectured at computer user's groups symposia, including SoftCon, EDGE and DECUS.

Web Development

Seybold Case Study and Technical Briefs

Chryste Sullivan
Channel Marketing Manager
Tools and Technology Group
Intel Corporation

Overview

Intel and the Intel® Architecture are gaining ever-increasing influence in the Web publishing arena. Showing its commitment to providing opportunities and exposure for Web tools and technologies companies, Intel presented one of the most popular and compelling 3D demos at Seybold 1999, San Francisco.

Using MetaFlash* and MetaStream* technology, jointly developed with MetaCreations Corporation, Intel captured 2D photographs of more than 1,200 Seybold attendees, converted them to 3D and posted them to a demo site. Using a PIN number, attendees could view and interact with their 3D faces and other models, send their PIN numbers to friends, and download their faces for use on their own Web sites. Even Web users who were unable to attend Seybold can experience 3D at this site. This demo is just one example of how Intel supports and shares successful tools and technologies with the industry.

Developer Pavillion

At its Developer and Workshop Pavilions, Intel introduced these 15 other leading-edge independent software vendors (ISVs), encompassing electronic and print publishing, streaming audio, and video, 3D design, and more.

- Adobe Systems
- Ideaworks3d
- Seer Systems
  - Apple Computer
  - Macromedia:
    Flash
    Dreamweaver*
- Shells Interactive
- Beatnik
- Digital Origin:
  PhotoDV*
  EditDV*
- Pulse Entertainment
- Eight Cylinder Studios
- RealNetworks
- Sonic Foundry
  Vegas Pro*
  Stream Anywhere*
  M.A.P. 2000*
- Equilibrium
- Right Hemisphere
  Deep Paint 3D*
  OneButton Productivity Solutions
- Smashing Ideas Animation

As Web site developers collaborate with these ISVs to implement cutting-edge tools and technology, they're finding new ways to differentiate themselves from their competition. They're taking full advantage of the power Intel Architecture brings to the Web.

Author Bio

Chryste Sullivan is a channel marketing manager in Intel's Tools and Technologies Group, where she focuses on the development community working with the Intel® Architecture. Chryste began her Intel career as a marketing communications manager and has been with the company for two years.

Prior to Intel, Chryste was a producer of children's multimedia titles for PrintPaks, now a division of Mattel. She also has experience selling print and CD ROM manufacturing for Quebecor, and began her career as a print production specialist for Microsoft. Chryste received her Bachelor of Science degree from the University of Oregon.

—End of Intel Developer Update Magazine Issue 1—